arXiv:1407.2109v2 [cs.DS] 5 May 2016

Planar Graphs: Random Walks and Bipartiteness Testing*

Artur Czumaj† Morteza Monemizadeh‡ Krzysztof Onak§ Christian Sohler¶

May 9, 2016

Abstract

We initiate the study of property testing in *arbitrary planar graphs*. We prove that *bipartiteness* can be tested in constant time: for every planar graph G and $\varepsilon > 0$, we can distinguish in constant time between the case that G is bipartite and the case that G is ε-far from bipartite. The previous bound for this class of graphs was $\tilde{O}(\sqrt{n})$, where n is the number of vertices, and the constant-time testability was only known for planar graphs with *bounded degree*. Our approach extends to arbitrary minor-free graphs.

Our algorithm is based on random walks. The challenge here is to analyze random walks for graphs that have good separators, i.e., bad expansion. Standard techniques that use a fast convergence of random walks to a uniform distribution do not work in this case. Informally, our approach is to self-reduce the problem of finding an odd-length cycle in a multigraph G induced by a collection of cycles to the same problem in another multigraph G' induced by a set of shorter odd-length cycles, in such a way that when a random walk finds a cycle in G' with probability $p > 0$, then it does so in G with probability $\lambda(p) > 0$. This reduction is applied until the cycles collapse to self-loops, in which case they can be easily detected.

While the analysis presented in this paper applies only to testing bipartiteness, we believe that the techniques developed will find applications to testing other properties in arbitrary planar (or minor-free) graphs, in a similar way as in the past the advances in testing bipartiteness led to the development of testing algorithms for more complex graph properties.

*A preliminary version of this paper appeared in *Proceedings of the 52th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 423–432, Palm Springs, CA, October 22–25, 2011. IEEE Computer Society Press, Los Alamitos, CA, 2011.

†Department of Computer Science, Centre for Discrete Mathematics and its Applications, University of Warwick. Email: A.Czumaj@warwick.ac.uk. Research supported in part by the *Centre for Discrete Mathematics and its Applications (DIMAP)*, EPSRC award EP/D063191/1, by EPSRC award EP/G064679/1, and by a Weizmann-UK Making Connections Grant "The Interplay between Algorithms and Randomness."

‡Department of Computer Science, Goethe-Universität Frankfurt, Germany and Department of Computer Science, University of Maryland at College Park, USA. Email: monemizadeh@em.uni-frankfurt.de and monemi@umiacs.umd.edu, Supported in part by MO 2200/1-1.

§IBM T.J. Watson Research Center. Email: konak@us.ibm.com. Supported in part by a Simons Postdoctoral Fellowship and NSF grants 0732334 and 0728645.

¶Department of Computer Science, TU Dortmund. Email: christian.sohler@tu-dortmund.de. Supported in part by the German Research Foundation (DFG), grant So 514/3-2 and ERC grant No. 307696.

1 Introduction

Property testing studies relaxed decision problems in which one wants to distinguish objects that have a given property from those that are far from this property (see, e.g., [6]). Informally, an object $\mathcal{X}$ is ε-far from a property $\mathcal{P}$ if one has to modify at least an ε-fraction of $\mathcal{X}$'s representation to obtain an object with property $\mathcal{P}$, where ε is typically a small constant. Given oracle access to the input object, a typical property tester achieves this goal by inspecting only a small fraction of the input. Property testing is motivated by the need to understand how to extract information efficiently from massive structured or semi-structured data sets using small random samples.

One of the main and most successful directions in property testing is *testing graph properties*, as introduced in papers of Goldreich et al. [7, 8]. There are two popular models for this task, which make different assumptions about how the input graph is represented and how it can be accessed.

For a long time, the main research focus has been on the *adjacency matrix model*, designed typically for *dense* graphs [7]. In this model, after a sequence of papers, it was shown that testability of a property in constant time is closely related to Szemerédi partitions of the graph. More precisely, a property is testable in constant time[1] if and only if it can be reduced to testing finitely many Szemerédi partitions [1].

The *adjacency list model* has been designed mostly for *sparse* graphs. In the most standard scenario, the model is studied with an additional restriction that the degree of the graph is at most a certain predefined constant d [8]. Unlike in the adjacency matrix model, it is not yet completely understood what graph properties are testable in constant time in the adjacency list model. Known examples include all hyperfinite properties [16] (see also [3] and [11] for previous general results), connectivity, k-edge-connectivity, the property of being Eulerian [8], and the property of having a perfect matching [17]. On the other hand, some properties testable in constant time in the dense graph model, such as bipartiteness and 3-colorability, are known to require a superconstant number of queries [4, 8].

Even less is known about efficiently testable properties for sparse graphs that do not have a degree bound. It turns out the constant degree bound in the adjacency list model is essential for many of the results mentioned above. All constant-time testers mentioned above use the fact that in a graph with constant maximum degree for a small sample set one can explore all vertices that have constant distance to one of the sample vertices and then decide based on the obtained information. It is known that connectivity, k-edge-connectivity, and Eulerian graphs are testable in constant time [15]. However, no general results characterizing constant-time testable properties are known.

Bipartiteness. The problem of testing bipartiteness has been a great benchmark of the capabilities of property testing algorithms in various graph models. It was one of the first problems studied in detail in both the dense graph model [7] and the bounded degree graph model [8, 9]. Bipartiteness is known to be testable in $\widetilde{O}(1/\varepsilon^2)$ time in the dense graph model [2]. However, in the bounded degree graph model, it requires $\Omega(\sqrt{n})$ queries [8] and is testable in $\widetilde{O}(\sqrt{n}\cdot\varepsilon^{-O(1)})$ time [9], where n is the number of vertices. Kaufman et al. [12] show that the property is still testable in $\widetilde{O}(\sqrt{n}\cdot\varepsilon^{-O(1)})$ time in the adjacency list model for graphs that have constant *average* degree.

Czumaj et al. [5] show in the bounded-degree model that if the underlying graph is planar, then any hereditary graph property[2], including bipartiteness, is testable in constant time. This approach can be generalized to any class of graphs that can be partitioned into constant-size components by removing εn edges

[1] Throughout the paper we say that a property is *testable in constant time* if there is a testing algorithm whose number of queries to the input is independent of the input size, possibly depending only on the proximity parameter ε.

[2] A graph property is *hereditary* if it is closed under vertex removals.

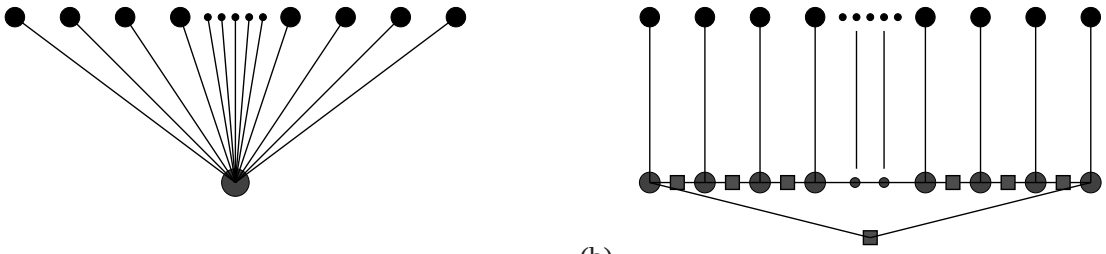

(a) (b)

Figure 1: An example of the process of splitting a vertex that reduces any graph into a graph of maximum degree at most 3 and that maintains planarity. For the graph in (a), Figure (b) depicts the splitting that is invariant to being bipartite.

of the graph, for any $\varepsilon > 0$. Graphs satisfying this property are called *hyperfinite*, and they include all bounded-degree minor-free graphs.

Hassidim et al. [11] show that in fact, the distance to most hereditary properties can be approximated in constant time in such graphs. These results are generalized in the work of Newman and Sohler [16], who show that in hyperfinite graphs, one can approximate the distance to *any graph property* in constant time. In particular, this implies that *any graph property is testable in constant time in hyperfinite graphs*, and therefore, in bounded-degree planar graphs.

The central goal of this paper is to initiate the research on the complexity of testing graph properties in general unbounded degree minor-free graphs. Our main technical contribution is the design and analysis of a constant time algorithm testing bipartiteness in arbitrary planar graphs. We show (in Theorem 4) that *a long enough constant-length random walk from a random vertex discovers an odd-length cycle in a graph far from bipartite with constant probability*. The result extends to an arbitrary family of minor-free graphs (Theorem 19 in Section 6).

1.1 Techniques

Our approach is based on a new analysis technique for random walks in planar (and minor-free) graphs. We first show that a planar graph that is far from bipartiteness has a linear number of edge-disjoint cycles of constant odd-length. Then we show a reduction that preserves up to a constant factor the probability of discovering an odd-length cycle by a random walk. We show that after a constant number of reductions we obtain a multigraph in which the probability of discovering an odd-length cycle is lower bounded by a positive constant.

1.2 Approaches that do not work

Given that bipartiteness can be tested in constant time in planar graphs of bounded degree [5], it may seem that there is a simple extension of this result to arbitrary degrees. We now describe two natural attempts at reducing our problem to testing bipartiteness in other classes of graphs. We explain why they fail. We hope that this justifies our belief that new techniques are necessary to address the problem.

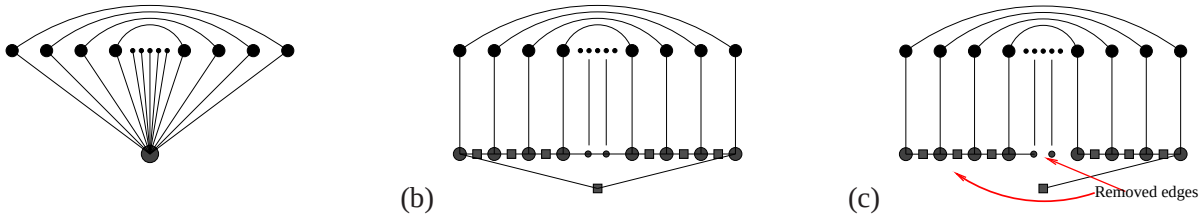

Figure 2: An example showing that the splitting construction from Figure 1 can reduce the distance from being bipartite. The planar graph in (a) (in which the i^{th} top vertex from the left is connected by an edge to the i^{th} top vertex from the right) has $\Theta(n)$ edge-disjoint cycles of length 3 and is ε-far from bipartite (one has to remove at least $\frac{n-1}{2}$ edges to obtain a bipartite graph). However, after the splitting, the obtained graph (Figure (b)) can be made bipartite just by removal of two edges: Figure (c) depicts a bipartite graph obtained after removal of such two edges: one of the two edges at the bottom and the middle edge in the split part.

1.2.1 Transforming into a constant-degree planar graph

The first and possibly the most natural approach to designing a constant-time algorithm for testing bipartiteness in arbitrary planar graphs would be to extend the known constant-time algorithm for *bounded-degree* planar graphs [5]. This could be achieved by first transforming an input planar graph G with an arbitrary maximum-degree into a planar graph $G^\star$ with bounded-degree and then running the tester for $G^\star$ to determine the property for G. However, we are not aware of a transformation that would behave well and we do not expect any such transformation to exist.

For example, one can reduce the maximum degree to at most 3 by splitting every vertex of degree $d > 3$ into d vertices of degree 3. It is also easy to ensure that this reduction maintains the planarity, and also the property of being bipartite (see Figure 1). However, there are two properties that are not maintained: one is the distance from being bipartite (see Figure 2) and another is that the access to the neighboring nodes requires more than constant time (though this can be "fixed" if one allows each vertex to have its adjacency list ordered consistently with some planar embedding). In particular, Figure 2 depicts an example of a planar graph that is originally ε-far from bipartite, but after the transformation it suffices to remove 2 edges to obtain a bipartite graph.

1.2.2 Substituting high-degree vertices with expanders

Another transformation of the graph is considered by Kaufman et al. [12]. They replace every high degree vertex with a constant-degree bipartite expander. While they prove that this construction preserves the distance, it is clear that it cannot preserve the planarity, since planar graphs are not expanders. However, for general graphs we know that testing bipartiteness requires $\Omega(\sqrt{n})$ queries [8] and we do not know how to exploit the structure of the graph after the transformation.

2 Preliminaries

Bipartiteness. A graph is *bipartite* if one can partition its vertex set into two sets A and B such that every edge has one endpoint in A and one endpoint in B. We also frequently use the well known fact that a graph is bipartite if and only if it has no odd-length cycle.

We now formally introduce the notion of being far from bipartiteness.[3] The notion is parameterized by a distance parameter $\varepsilon > 0$.

Definition 1 A graph $G = (V, E)$ is *ε-far from bipartite* if one has to delete more than $\varepsilon|V|$ edges from G to obtain a bipartite graph.

Property testing. We are interested in finding a *property testing algorithm* for bipartiteness in planar graphs, i.e., an algorithm that inspects only a very small part of the input graph, and accepts bipartite planar graphs with probability at least $\frac{2}{3}$, and rejects planar graphs that are ε-far away from bipartite with probability at least $\frac{2}{3}$, where ε is an additional parameter.

Our algorithm always accepts every bipartite graph. Such a property testing algorithm is said to have *one-sided error*.

Access model. The access to the graph is given by an *oracle*. We consider the oracle that allows two types of queries:

- *Degree queries:* For every vertex $v \in V$, one can query the degree of v.
- *Neighbor queries:* For every vertex $v \in V$, one can query its i^{th} neighbor.

Observe that by first querying the degree of a vertex, we can always ensure that the i^{th} neighbor of the vertex exists in the second type of query. In fact, in the algorithm that we describe in this paper, the neighbor query can be replaced with a weaker type of query: *random neighbor query*, which returns a random neighbor of a given vertex v; each time the neighbor is chosen independently and uniformly at random.

The *query complexity* of a property testing algorithm is the number of oracle queries it makes.

Basic properties of planar graphs. We extensively use the following well-known properties of planar graphs. The graph $G' = (V', E')$ obtained by the *contraction* of an edge $(u, v) \in E$ into vertex u is defined as follows: $V' = V \setminus \{v\}$ and $E' = \{(x, y) \in E : x \neq v \wedge y \neq v\} \cup \{(x, u) : (x, v) \in E \wedge x \neq u\}$. A graph G' that can be obtained from a graph G via a sequence of edge removals, vertex removals, and edge contractions is called a *minor* of G.

We use the following well-known property of planar graphs.

Fact 2 *Any minor of a planar graph is planar.*

Furthermore, we use the following upper bound on the number of edges in a simple planar graph, which follows immediately from Euler's formula.

Fact 3 *For any simple planar graph $G = (V, E)$ (with no self-loops or parallel edges), $|E| \leq 3|V| - 6$.*

We remark that for any class of graphs $\mathcal{H}$ that is defined by a finite collection of forbidden minors, similar statements are true, i.e., if $G \in \mathcal{H}$, then any minor of G also belongs to $\mathcal{H}$ and if $G = (V, E) \in \mathcal{H}$, then G has $O(|V|)$ edges (where the constant in the big O notation depends on the set of forbidden minors).

[3] The standard definition of being ε-far (see, for example, the definition in [12]) expresses the distance as the fraction of edges that must be modified in $G = (V, E)$ to obtain a bipartite graph. Compared to our Definition 1, instead of deleting $\varepsilon|V|$ edges, one can delete $\varepsilon|E|$ edges. For any class of graphs with an excluded minor, the number of edges in the graph is upper bounded by $C \cdot |V|$, where C is a constant. Moreover, unless the graph is very sparse (i.e., most of its vertices are isolated, in which case even finding a single edge in the graph may take a large amount of time), the number of edges in the graph is at least $\Omega(|V|)$. Thus, under the standard assumption that $|E| = \Omega(|V|)$, the ε in our definition and the ε in the previous definitions remain within a constant factor. We use our definition of being ε-far for simplicity; our analysis can be extended to the standard definition of being ε-far in a straightforward way.

Notation. Throughout the paper we use several constants depending on ε. We use lower case Greek letters to denote constants that are typically smaller than 1 (e.g., $\delta_i(\varepsilon)$) and lower case Latin letters to denote constants that are usually larger than 1 (e.g., $f_i(\varepsilon)$). All these constants are always positive. Furthermore, throughout the paper we use the asymptotic symbols $O_\varepsilon(\cdot)$, $\Omega_\varepsilon(\cdot)$, and $\Theta_\varepsilon(\cdot)$, which ignore multiplicative factors that depend only on ε and that are positive for $\varepsilon > 0$.

3 Algorithm Random-Bipartiteness-Exploration

We first describe our algorithm for testing bipartiteness of planar graphs with arbitrary degree and provide the high level structure of its analysis. Most of the technical details appear in Section 4.

Random-Bipartiteness-Exploration (G, ε):

- Repeat $f(\varepsilon)$ times:
  - Pick a random vertex $v \in V$.
  - Perform a random walk of length $g(\varepsilon)$ from v.
  - If the random walk found an odd-length cycle, then **reject**.
- If none of the random walks found an odd-length cycle, then **accept**.

Theorem 4 *There are positive functions f and g such that for every planar graph G, we have*

- *if G is bipartite, then* **Random-Bipartiteness-Exploration**(G, ε) *accepts G, and*
- *if G is ε-far from bipartite, then* **Random-Bipartiteness-Exploration**(G, ε) *rejects G with probability at least* 0.99.

We first observe that the first claim is obvious: if G is bipartite, then every cycle in G is of even length and hence **Random-Bipartiteness-Exploration** always accepts. Thus, to prove Theorem 4, it suffices to show that if G is ε-far from bipartite, then **Random-Bipartiteness-Exploration** rejects G with probability at least 0.99. Therefore, from now on, we assume that the input graph G is ε-far from bipartite for some constant $\varepsilon > 0$. Furthermore, note that it suffices to show that a *single* random walk of length $O_\varepsilon(1)$ finds an odd-length cycle with probability $\Omega_\varepsilon(1)$. Indeed, for any functions g and f, if a random walk of length $g(\varepsilon) = O_\varepsilon(1)$ finds an odd-length cycle with probability at least $5/f(\varepsilon) = \Omega_\varepsilon(1)$, then this implies that $f(\varepsilon) = O_\varepsilon(1)$ independent random walks detect at least one odd-length cycle with probability at least $1 - (1 - 5/f(\varepsilon))^{f(\varepsilon)} \geq 1 - e^{-5} \geq 0.99$. Therefore, in the remainder of the paper, we analyze algorithm **Random-Walk**(G, t) below. We have to prove that there is $t = t(\varepsilon) = O_\varepsilon(1)$ such that for every planar graph G that is ε-far from bipartite, **Random-Walk**(G, t) finds an odd-length cycle with probability $\Omega_\varepsilon(1)$. This implies Theorem 4.

Random-Walk (G, t):

- Pick a random vertex $v \in V$.
- Perform a random walk of length t from v.
- If the random walk found an odd-length cycle, then **reject**.
- If not, then **accept**.

3.1 Outline of the analysis of Random-Walk when G is ε-far from bipartite

Because of the arguments presented above, the remainder of the paper deals with the main technical challenge of our result: proving that our algorithm **Random-Walk** finds with sufficient probability an odd-length cycle in any planar graph $G = (V, E)$ that is ε-far from bipartite. To this end, we find a subgraph H of G that has the properties stated in the following lemma:

Lemma 5 *For every $\varepsilon \in (0, 1)$, there is a $t = t(\varepsilon) > 0$ such that for every graph $G = (V, E)$ that is ε-far from bipartite, there exists a subgraph H of G with the following properties:*

(a) if **Random-Walk***(H, t) finds an odd-length cycle in H with probability $\Omega_\varepsilon(1)$, then* **Random-Walk***(G, t) finds an odd-length cycle in G with probability $\Omega_\varepsilon(1)$, and*

(b) **Random-Walk***(H, t) finds an odd-length cycle in H with probability $\Omega_\varepsilon(1)$.*

If such a subgraph H always exists, these properties immediately imply that **Random-Walk**(G, t) finds an odd-length cycle in G with probability $\Omega_\varepsilon(1)$ and so by the discussion above, Theorem 4 follows.

In order to prove the existence of H, we construct a series of subgraphs $G \supseteq H_1 \supseteq \cdots \supseteq H_k$ for some $k = k(\varepsilon)$, such that each H_i, $1 \leq i \leq k$, satisfies property (a) and such that $H := H_k$ also satisfies property (b). We begin by constructing graph H_1, which consists of a union of $\Omega_\varepsilon(n)$ edge-disjoint short odd-length cycles from G. In order to construct H_1 we use the Klein-Plotin-Rao decomposition theorem, which shows that one can remove $\varepsilon|V|/2$ edges from G to partition G into connected components such that the every pair of vertices from the same component has distance $O_\varepsilon(1)$ in G. Then we show that there exists a component that contains a short odd-length cycle. We remove the cycle from the graph and repeat this process as long as we find a short odd-length cycle. The set of removed cycles forms a graph H_1', which after some further processing to satisfy property (a) is turned into H_1. The detailed construction appears in Section 4.1.

Then we design a reduction that takes a graph H_i consisting of a set C_i of $\Omega_\varepsilon(n)$ edge-disjoint short odd-length cycles and we construct from it a subgraph H_{i+1} that consists of a subset C_{i+1} of C_i of $\Omega_\varepsilon(n)$ cycles. To guide our process, we associate with each H_i a certain multigraph $\mathcal{H}_i$ that is constructed from H_i by performing edge contractions while keeping parallel edges and self-loops. The purpose of $\mathcal{H}_i$ is to approximate *how the random walk behaves at a larger scale*, i.e., one step of a (weighted) random walk in $\mathcal{H}_i$ corresponds to multiple steps in H_i. Each vertex u of $\mathcal{H}_i$ corresponds to a subset of vertices from V that contains u. Furthermore, these subsets are disjoint and form a partition of V. Vertex u from $\mathcal{H}_i$ can be thought of as a representative of the corresponding set. In particular, a random walk in $\mathcal{H}_i$ chooses its starting vertex u with probability proportional to the size of its represented subset. Edges in $\mathcal{H}_i$ represent paths in H_i; the paths represented by parallel edges may intersect. Parallel edges are also taken into account

when we perform a random walk. The probability to move from u to v is proportional to the number of parallel edges (u, v).

Let us recall that H_i consists of a set of short odd-length cycles C_i. The construction of $\mathcal{H}_i$ ensures that every cycle in C_i also corresponds to a cycle in $\mathcal{H}_i$. Clearly, the length of the corresponding cycle in $\mathcal{H}_i$ is at most the length of the cycle in H_i but due to edge contractions it may become lower. When we are constructing H_{i+1} from H_i by removing cycles from C_i to obtain C_{i+1}, we ensure that for every cycle from C_{i+1} the length in $\mathcal{H}_{i+1}$ is shorter than the length of the cycle in $\mathcal{H}_i$ (in addition to preserving property (a)). Thus, after k steps (when k is the maximum initial length of a cycle) all cycles are contracted to self-loops. But then, if $\mathcal{H}_k$ properly approximates the behavior of the random walk, a cycle can be detected and so $H = H_k$ satisfies both property (a) and (b), which finishes the proof.

While the above construction outlines the main line of thought, multiple details have to be taken care of. For instance, during the edge contractions we lose information about the original lengths of the paths and so an odd-length cycle in $\mathcal{H}_i$ may correspond to an even length set of cycles in H_i.

4 Analysis of Random-Walk when G is ε-far from bipartite

The main proof, our analysis of **Random-Walk** when G is ε-far from bipartite, can be kept relatively short, if we assume that the reduction from H_i to H_{i+1} works as supposed. Proving this reduction is the key technical challenge. We therefore first present the complete proof except for Lemma 11, whose proof appears in Section 5. Our proof follows the general outline sketched in the previous section. Our first step is to show how to obtain our starting subgraph H_1.

4.1 Finding the first subgraph H_1

We start by taking a closer look at property (a) of Lemma 5 and we show that it is implied by a simple condition on the degrees of the vertices in H, namely, the degree of each vertex is either 0 or no more than a constant factor smaller than its corresponding degree in G. Our construction of H_i in the remainder of the paper satisfies an alternate property (a'), defined below, which implies property (a).

Lemma 6 *Let $G = (V, E)$ be a graph, $t = \Theta_\varepsilon(1)$ be the length of a random walk, and H be a subgraph of G on vertex set V such that the following property holds:*

(a') for every vertex $v \in V$, either $deg_H(v) = 0$ or $deg_H(v) = \Omega_\varepsilon(deg_G(v))$.

Then property (a) of Lemma 5 is satisfied for a random walk of length t.

Proof. Consider a single walk w in H of length t that finds an odd-length cycle. Since H is a subgraph of G, the same walk exists in G. Furthermore, every vertex visited in w must have $\deg_H(v) = \Omega_\varepsilon(\deg_G(v))$. Therefore, at every step, the probability of following w decreases in G by at most a factor of $O_\varepsilon(1)$, compared to H. Overall the probability of w decreases by at most a factor of $(O_\varepsilon(1))^t = (O_\varepsilon(1))^{O_\varepsilon(1)} = O_\varepsilon(1)$. Summing up over all such walks proves the lemma. ■

We now proceed with the construction of H_1. We first construct a subgraph of G that is a union of $\Omega_\varepsilon(n)$ short edge-disjoint cycles. Then we modify it so that it satisfies property (a') and thus property (a).

We make use of the following Klein-Plotkin-Rao decomposition theorem [13].

Lemma 7 [13] *Let $G = (V, E)$ be a planar graph and let δ be a parameter, $0 < \delta < 1$. There is a set of at most $\delta|V|$ edges in G whose deletion decomposes G into connected components, where the distance (in the original graph G) between any two nodes in the same component is $O(1/\delta^2)$.*

Lemma 7 helps us realize the first part of the plan: showing that there are many short edge-disjoint odd-length cycles in a highly non-bipartite graph.

Lemma 8 (Many short odd-length cycles) *In every planar graph $G = (V, E)$ that is ε-far from bipartite, there exists a collection C of $\Omega_\varepsilon(|V|)$ edge-disjoint odd-length cycles of length at most $k = k(\varepsilon) = O(\varepsilon^{-2})$.*

Proof. We find the cycles one by one. Suppose that we have already found in G a set of ℓ edge-disjoint odd-length cycles of length at most $k = k(\varepsilon)$ each, where $\ell < \frac{\varepsilon|V|}{2k} = \Omega_\varepsilon(|V|)$. We show the existence of one more such cycle, which by induction yields the lemma. Let $G^\star$ be the subgraph of G obtained by removing the ℓ edge-disjoint odd-length cycles of length at most k each. Since $\ell < \frac{\varepsilon|V|}{2k}$, $G^\star$ is obtained by removing less than $\frac{\varepsilon|V|}{2}$ edges, and hence $G^\star$ is $\varepsilon/2$-far from bipartite. Apply Lemma 7 to $G^\star$ with $\delta = \frac{\varepsilon}{2}$ and let H be the resulting graph. Since $G^\star$ is $\varepsilon/2$-far from bipartite, H is not bipartite. Let us consider a connected component C_H of H that is not bipartite and let v be a vertex from C_H. Build a BFS tree from v in $G^\star$. Since C_H is not bipartite, there must exist two vertices u_1 and u_2 in C_H that have the same distance from v and that are connected by an edge in H (otherwise, we could define a bipartition of C_H by the parity of the distance from v in the BFS tree). Let v' be the last common vertex on the paths from v to u_1 and from v to u_2 in the BFS tree. The cycle in $G^\star$ that starts at v', goes to u_1 via the BFS tree edges, then takes the edge connecting u_1 and u_2, and finally returns to v' via the BFS tree edges. Let $k' = O(1/\varepsilon^2)$ be the bound on the diameter of C_H (in $G^\star$) that follows from Lemma 7. Since the BFS tree is a shortest path tree (from v), this cycle has length $k = k(\varepsilon) = 2k' + 1 = O(1/\varepsilon^2)$. ■

Given any set of cycles C on vertex set V, we write $G(C)$ to denote the graph on vertex set V that is induced by C, i.e., $G(C) = (V, E_C)$ with E_C being the union of the edges of the cycles in C. While Lemma 8 provides us with a graph $G(C)$ that has a linear number of disjoint short odd-length cycles, it is by no means clear that this new graph $G(C)$ satisfies property (a) of Lemma 5. However, we show in the next lemma that there is always a subset $C' \subseteq C$ with cardinality $|C'| = \Omega_\varepsilon(|C|)$ such that the graph $G(C')$ satisfies property (a) via showing that it satisfies property (a').

Lemma 9 (Transformation to obtain property (a')) *Let $G = (V, E)$ be a planar graph. Let C be a set of $\Omega_\varepsilon(|V|)$ edge-disjoint cycles on V in G, each of length at most k, for some $k = k(\varepsilon) = O_\varepsilon(1)$. Then there exists a subset $C' \subseteq C$ with $|C'| = \Omega_\varepsilon(|V|)$ such that the graph $G(C')$ satisfies condition (a') of Lemma 6. That is, for every $v \in V$, either $\deg_{G(C')}(v) = 0$ or $\deg_{G(C')}(v) = \Omega_\varepsilon(\deg_G(v))$.*

Proof. We construct the subset C' by deleting some cycles from C. The process of deleting the cycles is based on the comparison of the original degree of the vertices with the current degree in $G(C')$. To implement this scheme, we write $\deg_G(v)$ to denote the degree of v in the original graph G and we use the term *current degree* of a vertex v to denote its current degree in the graph $G(C')$ induced by the *current* set of cycles C' (where "current" means at a given moment in the process). Let $\alpha = |C|/|V| = \Omega_\varepsilon(1)$. We repeat the following procedure as long as possible: if there is a non-isolated vertex $v \in V$ with current degree in $G(C')$ at most $\frac{\alpha}{12}\deg_G(v)$, then we delete from C' all cycles going through v. To estimate the number of cycles deleted, we charge to v the number of deleted cycles in each such operation. Observe that each $v \in V$ will be processed not more than once. Indeed, once v has been used, it becomes isolated,

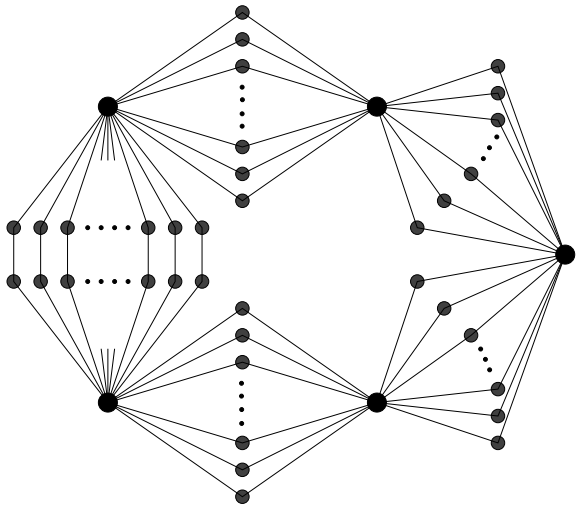

Figure 3: An illustrative example.

and hence it is not used again. Therefore, at most $\frac{\alpha}{12}\deg_G(v)$ cycles from C' can be charged to any single vertex. This, together with the inequality $\sum_{v\in V}\deg_G(v) \leq 6|V|$ by planarity of $G(C')$, implies that the total number of cycles removed from C to obtain C' is upper bounded by $\sum_{v\in V}\frac{\alpha}{12}\deg_G(v) \leq \frac{\alpha}{2}|V|$. Since $|C| = \alpha|V|$, we conclude that $|C'| \geq |C| - \frac{\alpha}{2}|V| = \frac{\alpha}{2}|V| = \Omega_\varepsilon(|V|)$. ■

Thus, by Lemma 9, we can construct a subgraph $H_1 := G(C')$ of G that is composed of a collection of $\Omega_\varepsilon(|V|)$ odd-length edge-disjoint cycles that satisfy property (a) of Lemma 5.

4.2 Constructing H_{i+1} from H_i

We begin with presenting some challenges of our construction and describe why two most natural lines of extending the analysis from Section 4.1 fail. After that, to facilitate our analysis, we first describe our framework in Section 4.2.2 and then, in Section 4.2.3, we present details of the construction of H_{i+1}.

4.2.1 On the challenges of the analysis

It is tempting to try to make a shortcut and avoid the need of constructing $H_2, \ldots, H_k$ and prove that a random walk finds an odd-length cycle in H_1 by showing that a fixed cycle is found with probability $\Omega_\varepsilon(1/|V|)$. Such a statement would be trivially true for graphs with a constant maximum degree, but it is false for arbitrary planar graphs, as illustrates the example in Figure 3. The graph in this example is composed of many parallel cycles that intersect in a few vertices of high degree. It is easy to see that a random walk finds an odd-length cycle with constant probability. However, any fixed cycle is only found with sub-constant probability.[4] This implies that in our arguments it is important to *exploit parallel-like structure* in the graphs.

Similarly, it could be tempting to hope that algorithm **Random-Walk** finds with constant probability an odd-length cycle in *an arbitrary graph* G (not necessarily planar) that is a union of a linear number of edge-disjoint odd-length cycles of constant length. This, however, turns out not to be the case. Consider an

[4]One can show that the probability of finding any fixed cycle is $n^{-\Omega_\varepsilon(1)}$.

expander with girth $\omega(1)$, in which every vertex has degree $\omega(1)$ (see, e.g., [14]). Then replace every edge by a cycle of length 3. It is easy to see that a random walk of constant length will find an odd-length cycle only if it finds one of the new cycles of length 3. However, since the probability of this event is $o(1)$, the algorithm will not find a cycle with constant probability. This implies that in our arguments it is crucial to *exploit planarity*.

4.2.2 Graph representation

All our graphs H_i are formed by sets C_i of short odd-length edge-disjoint cycles from $G = (V, E)$. We also have that $C_1 \supseteq C_2 \supseteq \cdots \supseteq C_k$ and $|C_k| = \Omega_\varepsilon(|V|)$, and thus $H_1 \supseteq H_2 \supseteq \cdots \supseteq H_k$. For a set of cycles C on vertex set V, we write $G(C)$ to denote the graph with vertex set V and edge set being the set of edges from the cycles in C. Thus, in our reduction $H_i = G(C_i)$. Recall that in addition to the graphs H_i we use the multigraphs $\mathcal{H}_i$. In order to define $\mathcal{H}_i$, we first require a partition $P_i : V \to V$ of the vertices that describes how the edges are contracted. The idea is that the vertex set $P^{-1}(u)$ is contracted into vertex u. In order to be meaningful, this partition has to satisfy the properties given in the following definition.

Definition 10 Let C be an arbitrary set of disjoint cycles in G. A partition $P : V \to V$ is called *good* for C if it satisfies the following four properties.

- If $u \in V$ is in the image of P, then $P(u) = u$.
- If $u \in V$ is not contained in any cycle from C, then $P(u) = u$.
- For each cycle $\mathfrak{c} \in C$ and each partition class $P^{-1}(u)$, where u belongs to the image of P, if $P^{-1}(u)$ contains a vertex of $\mathfrak{c}$, then $\mathfrak{c}$ also contains u.
- For each cycle $\mathfrak{c} \in C$ and each partition class $P^{-1}(u)$, where u belongs to the image of P, one of the following is true: *(i)* $P^{-1}(u)$ contains all vertices of $\mathfrak{c}$, or *(ii)* $P^{-1}(u)$ contains no vertex of $\mathfrak{c}$, or *(iii)* $P^{-1}(u)$ induces a path in $\mathfrak{c}$.

A vertex u in the image of P is called *the head of the partition class* $P^{-1}(u)$.

An example of a good partition is presented in Figure 4 (left side). The first property ensures that all partition classes have a proper head, i.e., a vertex into which all other vertices are contracted. The second property ensures that each isolated vertex in $G(C)$ has its own partition class. The third and fourth properties ensure that we can apply the contractions to each cycle from C individually. This way we may also obtain from C a multiset of contracted cycles $\mathcal{C}$.

Notice that the above definition implies that any partition class $P^{-1}(u)$, u in the image of P, forms a connected subgraph in $G(C)$. Indeed, assume that v and w are in $P^{-1}(u)$, $v \neq w$. By the second property, any of v and w which is not equal to u is contained in a cycle. By the third property, either $v = u$ or v is contained in a cycle, and hence v and u are in the same cycle; similarly, either $w = u$ or w and u are in the same cycle. Therefore, there is a path connecting v and w.

In our analysis, we modify a given partition P in two ways: by *edge contractions* and by *cycle removals*. As a result of contracting an edge (u, v), the partition classes of u and v are merged. Deleting one (or more) cycles from C may create an isolated vertex (with degree 0) in $G(C)$. In this case, the second property in Definition 10 may be violated. Therefore, whenever we create an isolated vertex u in $G(C)$ by deleting a cycle from C, we define $P(u) = u$ to satisfy the second condition. It is easy to verify that the other conditions are still satisfied after this modification has been applied to all newly created isolated vertices.

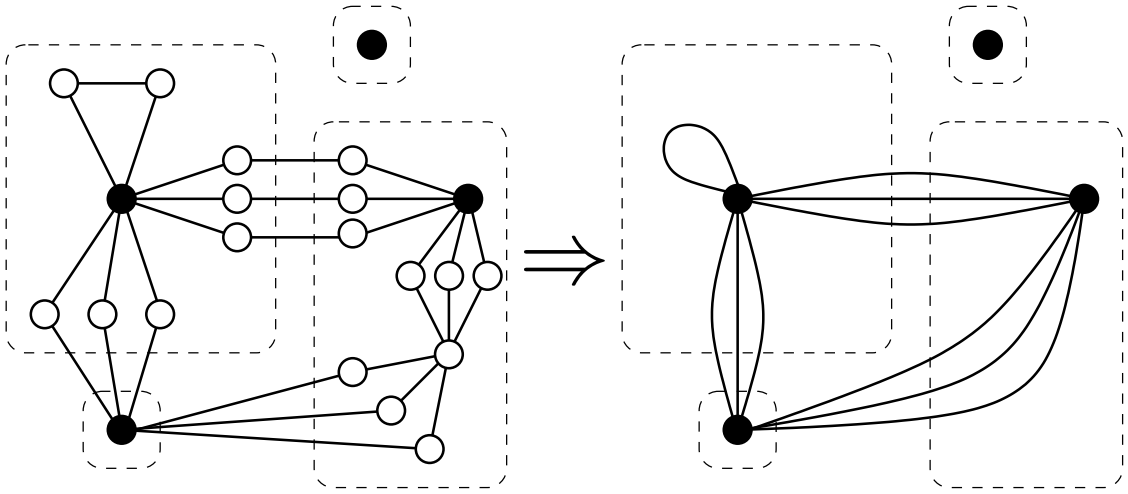

Figure 4: A sample good partition of vertices and the resulting contracted graph. The solid circles denote heads of partition classes.

We abuse the notation and write $P(\mathfrak{c})$ to denote the cycle obtained from $\mathfrak{c}$ by contracting its vertices according to P. We remark that the contracted cycles may be degenerated to self-loops or cycles of the form (u, v, u). For example, if all vertices of $\mathfrak{c}$ are contained in the same partition class of P, then the resulting cycle is a self-loop incident to the head of the partition class. For any set of cycles C, we also write $P(C)$ to denote the union of $P(\mathfrak{c})$ for all $\mathfrak{c} \in C$. Finally, we define $\mathcal{G}_P(C)$ to be the multigraph whose vertex set is the image of P and whose edge multiset is the multiset of edges of cycles in $P(C)$. The right side of Figure 4 contains the resulting multigraph for the input set of cycles and partition on the right. For a multigraph $\mathcal{G}$, we write $\deg_{\mathcal{G}}(u)$ to denote the number of edges in $\mathcal{G}$ that are incident to u (i.e., self-loops contribute 1 to the degree).

In our main reduction, we define $\mathcal{H}_i = \mathcal{G}_{P_i}(C_i)$. The multigraph $\mathcal{H}_i$ has the following interpretation. An edge (u, v) in $\mathcal{H}_i$ represents a path of H_i of length $O_\varepsilon(1)$. Therefore, with probability $\Theta_\varepsilon(1/\deg_{\mathcal{H}_i}(u))$, a random walk on $G(C_i)$ of (suitable) constant length starting in u reaches v. If there are q edges (u, v) in $\mathcal{H}_i$ then their union represents a subgraph of H_i that for random walks behaves like a set of q parallel non-intersecting paths. In particular, the probability of moving from u to v is $\Theta_\varepsilon(q/\deg_{\mathcal{H}_i}(u))$, where q is the number of edges (u, v) in $\mathcal{H}_i$.

Furthermore, a random walk starting from a vertex in $P^{-1}(u)$ reaches vertex u with constant probability after a constant number of steps.

With this definition, we can extend our notion of random walks to multigraphs in the following natural way (where we define the notion of odd-parity cycles below).

Random-Walk $(\mathcal{G} = (U, \mathcal{E}), t)$:

- Pick a random vertex $u \in U$, such that any $u \in U$ is chosen with probability $\frac{|P^{-1}(u)|}{|V|}$.
- Perform a random walk of length t from v, where the probability to move from a vertex $u \in U$ to a vertex $w \in U$ is $|\{(u, w) \in \mathcal{E}\}|/\deg_{\mathcal{G}}(u)$.
- If the random walk found an odd-parity cycle, then **reject**.
- If not, then **accept**.

Parities of edges and lengths of cycles. Since the contraction of edges influences the length of cycles and since we are searching for cycles of odd-length in $G(C)$, we need a way to keep track of the parity of the lengths of the cycles and paths explored and contracted. It would be easy, if we were just dealing with fixed cycles, but since we would like to allow combinations of cycles, we need to encode the contractions in the graph in a way that allows us to use combinations. This can be done as follows. For every cycle $\mathfrak{c} \in C_i$, we have a corresponding contracted cycle in $\mathfrak{c}' \in P_i(C_i)$ on vertex set U. Each edge in $\mathfrak{c}'$ corresponds to a path in $\mathfrak{c}$ and we define the parity of an edge (u, v) in $\mathfrak{c}'$ as 0, if the length of the path from u to v in $\mathfrak{c}$ is *even* and 1, if the length is *odd*. In the case that u and v are the only two vertices in $\mathfrak{c}'$, there are two edges connecting u and v and one of them has parity 0 while the other one has parity 1. This way, a cycle in the multigraph $\mathcal{G}$ has odd parity if it contains an odd number of edges with parity 1.

4.2.3 Main reduction from H_i and $\mathcal{H}_i$ to H_{i+1} and $\mathcal{H}_{i+1}$

The next lemma states our main technical contribution: the main reduction with its properties. We apply it several times to reduce the lengths of our cycles to 1. For the clarity of presentation, we postpone the proof of this lemma to Section 5.

Lemma 11 *Let $t = O_\varepsilon(1)$, let $G = (V, E)$ be a planar graph, and let $\ell \in \mathbb{N}$, $\ell \geq 2$. Let C be a set of $\Omega_\varepsilon(|V|)$ odd-length cycles in G. Let P be a partition that is good for C such that all cycles in $P(C)$ have length at most ℓ. Then we can construct a set of cycles $C^* \subseteq C$ with $|C^*| = \Omega_\varepsilon(|C|)$, and a partition P^* that is good for C^*, such that the following properties are satisfied:*

- *every cycle in $P^*(C^*)$ has length at most $\ell - 1$, and*
- *if the probability that* **Random-Walk**$(\mathcal{G}_{P^*}(C^*), t)$ *finds an odd-parity cycle is $\Omega_\varepsilon(1)$, then also the probability that* **Random-Walk**$(\mathcal{G}_P(C), 3t)$ *finds an odd-parity cycle is $\Omega_\varepsilon(1)$.*

This lemma is used to construct H_{i+1} and $\mathcal{H}_{i+1}$. We take as the input $C = C_i$ and $P = P_i$, and apply Lemma 11 to obtain $C_{i+1} = C^*$ and $P_{i+1} = P^*$, giving $H_{i+1} = G(C^*)$ and $\mathcal{H}_{i+1} = \mathcal{G}_{P^*}(C^*)$.

4.2.4 Useful property of $\mathcal{H}_k$

The first property of the construction in Lemma 11 implies that the length of the cycles in $P_i(C_i)$ decreases with increasing i. We therefore apply Lemma 11 $k - 1$ times, where k is the original upper bound for the cycle length in Lemma 9. As a result, we obtain a sequence $C_1 \supseteq \cdots \supseteq C_k$ consisting of sets of cycles

such that $P_k(C_k)$ consists solely of self-loops. We use the following property of graphs induced by such self-loops.

Lemma 12 *Let $G = (V, E)$ be a planar graph. Let C_k be a set of edge disjoint odd-length cycles in G and P_k be a partition that is good for C_k. If all cycles in $P_k(C_k)$ are self-loops and $|C_k| = \Omega_\varepsilon(|V|)$, then the probability that a 1-step random walk finds an odd-parity cycle in $\mathcal{H}_k$ is $\Omega_\varepsilon(1)$.*

Proof. Since C_k is a collection of $\Omega_\varepsilon(|V|)$ edge disjoint cycles in G, the underlying graph H_k has $\Omega_\varepsilon(|V|)$ edges. Therefore since H_k is planar, H_k must have $\Omega_\varepsilon(|V|)$ non-isolated vertices. Each of these non-isolated vertices is contracted in P_k into a vertex that is incident to a self-loop. Therefore, the probability to sample a vertex incident to a self-loop is $\Omega_\varepsilon(1)$ and then in one step we find a self-loop. By the definition of parities, the detected self-loop has odd parity. ■

4.3 Proof of Theorem 4: analyzing Random-Bipartiteness-Exploration

Now we are ready to prove our main theorem, Theorem 4. We already observed that it suffices to prove that for any planar graph $G = (V, E)$ that is ε-far from bipartite, **Random-Walk**(G, t) finds an odd-length cycle with probability $\Omega_\varepsilon(1)$ for $t = O_\varepsilon(1)$. We know by Lemma 8 that G contains a set C of $\Omega_\varepsilon(|V|)$ odd-length cycles of length at most $k = O_\varepsilon(1)$. Furthermore, by Lemma 9, there is a set $C_1 \subseteq C$ of $\Omega_\varepsilon(|V|)$ odd-length cycles of length at most k such that if **Random-Walk**$(G(C_1), t)$ on $G(C_1)$ finds an odd-length cycle with probability $\Omega_\varepsilon(1)$, then so does **Random-Walk**(G, t).

We then define $H_1 = G(C_1)$ and $P_1 : V \to V$ to be the identity. Clearly, P_1 is good for C_1. Then we apply Lemma 11 $k-1$ times to obtain sets of cycles C_i and partitions P_i that satisfy the properties of Lemma 11. In particular, since $k = O_\varepsilon(1)$, we know that $\mathcal{H}_k$ contains $\Omega_\varepsilon(|V|)$ self-loops. Moreover, if **Random-Walk**$(\mathcal{H}_k, T)$ finds an odd-parity cycle with probability $\Omega_\varepsilon(1)$, then so does **Random-Walk**$(\mathcal{H}_1, T \cdot 3^{k-1})$, which in turn implies that **Random-Walk**$(G, T \cdot 3^{k-1})$ also finds an odd-length cycle with probability $\Omega_\varepsilon(1)$. Thus, we only need to prove that **Random-Walk**$(\mathcal{H}_k, T)$ with $T = O_\varepsilon(1)$ finds an odd-parity cycle with probability $\Omega_\varepsilon(1)$. This follows immediately from Lemma 12.

5 Proof of Lemma 11

To complete the analysis, it remains to prove Lemma 11. We start with an overview of the proof. The main idea is to thin out the current set of cycles C to ensure that we can define a set of "contractions" to decrease the length of each of the remaining cycle. Further care is needed to ensure that after performing the contractions, we still maintain a good partition for the set of remaining cycles. This means that we are not allowed to contract edges that "shortcut" other remaining cycles. In order to avoid this we ensure that one of the contracted vertices, say v, either has only one distinct neighbor in $\mathcal{H}_i$ (in which case all cycles that involve this vertex must be of the form (u, v, u)) or it has exactly two distinct neighbors, say x and y, and all cycles that involve v contain the edges (x, v) and (v, y). Furthermore, in the latter case, we ensure that all edges (x, v) have the same parity and all edges (v, y) have the same parity, as otherwise the parity after the contraction would not be well-defined. This is captured by the following definition.

Definition 13 Let C be a set of cycles on vertex set V and P be a partition that is good for C. We say a vertex $v \in P(V)$ is *well-contractible* in $\mathcal{G}_P(C)$ if it satisfies one of the following conditions:

- v has only one distinct neighbor in $\mathcal{G}_P(C)$, or

- v has two distinct neighbors x and y in $\mathcal{G}_P(C)$, and
  - $\diamond$ all cycles in C that contain v contain also both x and y, and
  - $\diamond$ all edges (x, v) have the same parity and all edges (v, y) have the same parity.

Using this definition, our goal is to thin out the cycles such that every cycle has a well-contractible vertex. We also need to ensure that our contractions using well-contractible vertices do not interfere with each other. Therefore we additionally require that the set of well-contractible vertices forms an independent set.

After finding a set of cycles C' (and good partition P') with the above properties, we still have to do a "clean-up" phase to ensure that if a random walk finds with constant probability a cycle in $\mathcal{G}'_P(C')$, then this is also true in $\mathcal{G}_P(C)$. In order to achieve this, we use Lemma 9 to ensure that the degree of every non-isolated vertex in $\mathcal{G}'_P(C')$ is at least a constant fraction of its degree in $\mathcal{G}_P(C)$.

Outline. The roadmap is now as follows. We first find a subset of cycles such that every cycle has a well-contractible vertex. In order to do so, we first construct a large subset of cycles such that every cycle has a vertex with at most 6 distinct neighbors (Lemma 14). Then we argue that such a set of cycles satisfies at least one of the following two conditions: (1) It already has many self-loops, in which case we can just take this subset of self-loops and we are done with our reduction, or (2) we can remove all self-loops and process the remaining (linear number of) cycles to ensure that every cycle contains a well-contractible vertex (Lemma 17).

Lemma 14 *Let C be a set of edge-disjoint cycles on vertex set V that are of length at most k and such that $G(C)$ is a planar graph. Let P be a partition that is good for C. Then there is a set $C' \subseteq C$ of size at least $\frac{1}{4k+2}|C|$ such that every cycle in $P(C')$ has a vertex with at most 6 distinct neighbors in $\mathcal{G}_P(C')$.* [5]

Proof. We prove the lemma by presenting an algorithm that takes as its input a set of cycles C with planar $G(C)$ and a partition P that is good for C. The algorithm computes a subset C' that satisfies the properties of the lemma. The algorithm consists of two phases.

In the first phase, we partition C into *levels*, iteratively removing the cycles until C is empty. We start with C being the input set of cycles. In the j^{th} iteration, we choose an arbitrary vertex u_j with at most 6 distinct neighbors in $\mathcal{G}_P(C)$ and at least one incident cycle. Such a vertex exists since $G(C)$ is a planar graph, and so by Euler's formula there exists a vertex in $\mathcal{G}_P(C)$ with at most 5 distinct neighbors other than itself. Taking into account that $\mathcal{G}_P(C)$ may also contain self-loops, we find a vertex with at most 6 distinct neighbors. Here C refers to the current set C, i.e., after the removal of the sets from the previous iterations of the repeat-loop. Every cycle from C that contains u_j is removed from C. If a cycle $\mathfrak{c}$ is removed in the j^{th} iteration, its *level* $\ell(\mathfrak{c})$ is defined to be j.

In the second phase, we start again with C being the input set of cycles. We now iterate through the levels in decreasing order. For each level j, we let $A(j)$ denote the current subset of cycles in C at level j. Note that by definition of the level, all cycles in $A(j)$ must contain vertex u_j. Furthermore, we define $B(j)$ to be the subset of cycles of C that contain u_j and have a level smaller than j. We observe that if we remove all cycles in $B(j)$ from C, then every cycle in $A(j)$ contains a vertex (e.g., vertex u_j) with at most 6 distinct neighbors in $\mathcal{G}_P(C)$. The second phase explores this observation. For each j, we compare the size of $A(j)$ to the size of $B(j)$. If $A(j)$ is sufficiently large, i.e., at least a $\frac{1}{2k}$ fraction of $B(j)$, then we keep

[5] We slightly abuse notation here as P may be no longer a good partition due to vertex removal. In this case, we can still define $\mathcal{G}_P(C')$ the same way as above.

$A(j)$ and remove $B(j)$ from C; otherwise, we remove $A(j)$. Below we argue that at most half of the cycles from C are removed because they are in some removed set $A(j)$. Furthermore, for every $2k$ cycles that are removed because they are contained in a set $B(j)$, at least one cycle remains in C. This allows us to deduce the lemma. Details follow after the pseudocode describing more formally the process.

Assigning-Levels (set C of cycles and a partition P that is good for C)

- $j = 1$
- $C'' = C$

Phase 1:

- Repeat until C is empty:
  - Let u_j be a non-isolated vertex that has at most 6 distinct neighbors in $\mathcal{G}_P(C)$
  - For all cycles $\mathfrak{c} \in C$ that contain u_j, let $\ell(\mathfrak{c}) = j$
  - Remove from C all cycles that contain u_j
  - $j = j + 1$

Phase 2:

- $C = C''$
- Repeat until $j = 1$:
  - $j = j - 1$
  - $A(j) = \{\mathfrak{c} \in C : \ell(\mathfrak{c}) = j\}$
  - $B(j) = \{\mathfrak{c} \in C : \ell(\mathfrak{c}) < j$ and $\mathfrak{c}$ contains $u_j\}$
  - if $|A(j)| \geq \frac{1}{2k} \cdot |B(j)|$ then $C = C \setminus B(j)$ else $C = C \setminus A(j)$
- Return C

It remains to prove the correctness of the algorithm. We first observe that Phase 1 terminates since $\mathcal{G}_P(C)$ is a planar graph and therefore by Euler's formula, it has a vertex with at most 6 neighbors (this also holds during the execution of the algorithm since planarity is closed under edge removal and contractions).

It remains to analyze Phase 2 of the algorithm. Every cycle in C (a) is removed because it is contained in some set $A(j)$ that is removed from C in Phase 2, or (b) is removed because it is contained in some set $B(j)$ that is removed from C in Phase 2, or (c) is not removed and stays in the final set C. Let a, b, and c be the respective numbers of cycles. Clearly, $|C| = a + b + c$ and to prove the lemma we have to show that $c \geq |C|/k$. We proceed in two steps. We first prove that $a \leq \frac{1}{2}|C|$, which implies that $b + c \geq \frac{1}{2}|C|$. Then we argue that $2kc \geq b$. This yields $(2k + 1)c \geq \frac{1}{2}|C|$ and hence $c \geq \frac{1}{4k+2} \cdot |C|$.

Claim 15 $a \leq \frac{1}{2}|C|$.

Proof. We charge the vertices from the removed sets $A(j)$ to the sets $B(j)$ and derive a bound on the sum of sizes of the sets $B(j)$. Recall that every cycle contains at most k vertices. In

every cycle, one vertex is the vertex that has degree at most 6 (in $\mathcal{G}_P(C)$) when the cycle is removed in Phase 1 of the algorithm. Thus, every cycle is contained in at most $k-1$ different sets $B(j)$. It follows that

$$\sum_j |B(j)| \leq (k-1) \cdot |C|.$$

Let R denote the set of indices j such that $A(j)$ is removed from C during Phase 2. Observe that whenever we remove a set $A(j)$, we have $|A(j)| < \frac{1}{2k}|B(j)|$ by the condition in the process. It follows that

$$a = \sum_{j \in R} |A(j)| < \sum_{j \in R} \frac{1}{2k}|B(j)| \leq \frac{k-1}{2k} \cdot |C| < \frac{1}{2} \cdot |C|.$$

■

Claim 16 $2kc \geq b$.

Proof. For every set $B(j)$ removed from C, we know that $|A(j)| \geq \frac{1}{2k}|B(j)|$. At the point of time when $B(j)$ is removed from C, the set $A(j)$ remains in C because $A(j)$ and $B(j)$ are disjoint. Since we are iterating downwards through the levels of the cycles, the set $A(j)$ is also disjoint from all sets $B(j')$, $j' < j$, and so it is not removed also in any future iteration of the repeat loop. Thus, in this case each cycle from $A(j)$ remains in C until the end of the process and contributes to the value of c. Let R' be the set of indices j such that $A(j)$ remains in C during Phase 2 (and hence $B(j)$ is removed from C). Since each cycle of $A(j)$, $j \in R'$, contributes to c and since sets $A(j)$ are disjoint, we obtain $\sum_{j \in R'} |A(j)| \leq c$. Hence,

$$\frac{1}{2k}b = \frac{1}{2k} \sum_{j \in R'} |B(j)| \leq \sum_{j \in R'} |A(j)| \leq c,$$

which implies the claim. ■

This finishes the proof of Lemma 14, which follows from Claims 15 and 16 as argued above. ■

Our next lemma shows that if there are no self-loops, then we can cover a large number of cycles by well-contractible vertices.

Lemma 17 *Let C be a set of edge disjoint cycles on a vertex set V of length at most k. Let P be a partition that is good for C such that $\mathcal{G}_P(C)$ contains no self-loops. Let Q be the set of vertices in V that have at most 6 distinct neighbors in $\mathcal{G}_P(C)$. If every cycle in C contains at least one vertex from Q, then there is a subset $C' \subseteq C$, $|C'| \geq 12^{-2k}|C|$, such that every cycle $\mathfrak{c} \in C'$ has a vertex $v \in Q$ that is a well-contractible vertex in $\mathcal{G}_P(C')$.*[6]

Proof. For each vertex $v \in Q$, we select x and y independently uniformly at random with among its neighbors in $\mathcal{G}_P(C)$. If $x = y$, we delete from C all cycles $\mathfrak{c}$ that contain v and for which $P(\mathfrak{c})$ contains any vertex other than v or x. If $x \neq y$, we also select independently uniformly at random parities $p_x, p_y \in \{\text{odd}, \text{even}\}$. In this case, we remove from C every cycle $\mathfrak{c}$ that contains v unless $P(\mathfrak{c})$ contains both the

[6]We slightly abuse notation here as P may be no longer a good partition due to vertex removal. In this case, we can still define $\mathcal{G}_P(C')$ the same way as above.

edge (v,x) of parity p_x and the edge (v,y) of parity p_y. Let C' be the set of remaining cycles. We observe that every cycle $\mathfrak{c}$ in C' contains a vertex $v \in Q$ that is well-contractible (in fact, every vertex from Q contained in $\mathfrak{c}$ is well-contractible). The probability that a fixed cycle $\mathfrak{c}$ is not deleted is at least 12^{-2k}. Thus, the expected size of C' is at least $12^{-2k}|C|$, and therefore, there exists a set C' of that size that satisfies the lemma. ■

We now proceed to the main technical lemma that prepares our reduction step. Given a set C of cycles and a partition P, we show that there is a set Q and a subset C' of C of size $\Omega_\varepsilon(|C|)$ such that we can simultaneously contract edges at all vertices from Q to shorten all cycles in $P(C')$. Furthermore, the degree of each non-isolated vertex in the multigraph $\mathcal{G}_P(C')$ is comparable to its degree in $\mathcal{G}_P(C)$.

Lemma 18 *Let $\varepsilon \in (0,1)$. Let C be a set of edge-disjoint cycles on vertex set V of length at most $k = O_\varepsilon(1)$ such that $G(C)$ is planar and let P be a partition that is good for C. There exists a set $C' \subseteq C$ of size $\Omega_\varepsilon(|C|)$, a set of vertices $Q \subseteq P(V)$, and a partition P' that is good for C', such that*

- *P' is obtained from P by setting for every vertex u, $P'(u) = \begin{cases} P(u) & \text{if } \deg_{G(C')}(u) > 0, \\ u & \text{otherwise,} \end{cases}$*
- *Q is an independent set in $\mathcal{G}_{P'}(C')$,*
- *every vertex in Q is well-contractible in $\mathcal{G}_{P'}(C')$,*
- *every cycle in C' that is not a self-loop contains a vertex from Q, and*
- *every non-isolated vertex v from $\mathcal{G}_{P'}(C')$ satisfies $\deg_{\mathcal{G}_{P'}(C')}(v) = \Omega_\varepsilon(\deg_{\mathcal{G}_P(C)}(v))$.*

Proof. Let $C_1 \subseteq C$ be the subset of cycles from C that are self-loops in $P(C)$ and let $C_2 = C \setminus C_1$. If $|C_1| \geq |C_2|$, then we choose $C' = C_1$, $Q = \emptyset$ and P' as in the lemma to satisfy the properties of Lemma 18.

It remains to consider the case that $|C_2| > |C_1|$, in which case we select $C' \subseteq C_2$. We first set $P' = P$ and then iteratively modify P' such that it remains a good partition for the current set of cycles and can be obtained as the statement of the lemma specifies.

We modify P' in the way described earlier in the case of deletions of cycles. First, we set $P'(u) = u$ for every vertex u with degree 0 in $G(C_2)$. Then we apply Lemma 14 with C_2 and P' to obtain a set of cycles C_2' of size $\Omega_\varepsilon(|C|)$, such that every cycle from $P'(C_2')$ contains at least one vertex with at most 6 distinct neighbors. In order to maintain the first property in the lemma statement, we modify P' by setting $P'(v) = v$ for new isolated vertices in $G(C_2')$. Note that after the modification every cycle from $P'(C_2')$ still contains at least one vertex with at most 6 distinct neighbors.

Next, we apply Lemma 17 to C_2' and the current P'. We obtain sets Q' and $C_2'' \subseteq C_2'$. Q' is the set of vertices that have at most 6 distinct neighbors in $\mathcal{G}_{P'}(C_2'')$. $|C_2''| = \Omega_\varepsilon(|C_2'|) = \Omega_\varepsilon(|C|)$ and every cycle in C_2'' contains a vertex from Q' that is well-contractible. Then we apply Lemma 9 to C_2'' to obtain the final set of cycles $C' = \Omega_\varepsilon(|C_2''|)$. We modify P' as before, by setting $P'(v)$ for new isolated vertices in $\mathcal{G}_{P'}(C')$. Clearly, C' and P' satisfy the first and the fifth property specified in the lemma.

We set Q to be any maximal independent subset of vertices in Q' that are well-contractible in $\mathcal{G}_{P'}(C')$. This ensures both the second and third property. It remains to show that the fourth property is satisfied. By the construction of C', every cycle $\mathfrak{c} \in C'$ contains a vertex $v \in Q'$ that is well-contractible in $\mathcal{G}_{P'}(C')$. If $v \in Q$, we are done. Otherwise, since Q is maximal, we know that a neighbor of v in $\mathcal{G}_{P'}(C')$ is in Q. By the definition of well-contractible vertices, this neighbor is also in $\mathfrak{c}$, which completes the proof of the lemma. ■

5.1 Finalizing the proof of Lemma 11

We are ready to complete the proof of Lemma 11. For a planar graph $G = (V, E)$, let C be a set of $\Omega_\varepsilon(|V|)$ odd-length cycles in G and let P be a partition that is good for C such that all cycles in $P(C)$ have length at most $\ell \geq 2$. We first apply Lemma 18 to obtain the sets C' and Q and the partition P'. We set $C^* = C'$.

Now we consider all vertices in Q and contract each of them into one of its neighbors. This construction is well-defined since Q is an independent set by the second property of Lemma 18. For each vertex $u \in Q$, we select one of its neighbors in $\mathcal{G}_{P'}(C^*)$ and call it $\gamma(u)$. We want to contract all edges $(u, \gamma(u))$. A refinement P^* of P' is created as follows. For each vertex $u \in Q$ and for each vertex $v \in P'^{-1}(u)$, we set $P^*(v) = \gamma(u)$. For all other vertices $v \in V$, we set $P^*(v) = P'(v)$. One can easily verify that the resulting mapping P^* is a good partition for C^*.

Our next step is to argue that the length of the cycles in $P^*(C^*)$ is at most $\ell - 1$. We first note that the fourth property in the statement of Lemma 18 ensures that every cycle $\mathfrak{c} \in C^*$ that is not a self-loop contains a vertex $u \in Q$. This vertex is well-contractible in $\mathcal{G}_{P'}(C^\star)$—due to the third property—and our construction above contracts an edge in $\mathfrak{c}$ incident to this vertex.

Therefore, to complete the proof, we only need to argue that if **Random-Walk**$(\mathcal{G}_{P^*}(C^*), t)$ finds an odd-parity cycle with probability $\Omega_\varepsilon(1)$, then the probability that **Random-Walk**$(\mathcal{G}_P(C), 3t)$ finds an odd-parity cycle is also $\Omega_\varepsilon(1)$.

Each edge (u, v) in $\mathcal{G}_{P^*}(C^*)$ belongs to a cycle $\mathfrak{c}$ in $P^*(C^*)$. By our construction, a given edge (u, v) was either already present in $P(\mathfrak{c})$, or is a result of a contraction and there is a vertex $q \in Q$ such that (u, q) and (q, v) are edges of $P(\mathfrak{c})$. Note that in either case, the parity of the corresponding path from u to v is maintained. The probability to move from u to v via edge (u, v) in $\mathcal{G}_{P^*}(C^*)$ is $1/\deg_{\mathcal{G}_{P^*}(C^*)}(u)$. If (u, v) was already present in $\mathcal{G}_P(C)$, then the probability to move from u to v via edge (u, v) in $\mathcal{G}_P(C)$ is $1/\deg_{\mathcal{G}_P(C)}(u)$. The case that (u, v) corresponds to (u, q) and (q, v) in $\mathcal{G}_P(C)$ is more complicated since parallel edges become relevant. Since q is a well-contractible vertex in $\mathcal{G}_P(C')$, we know that all cycles in $P(C')$ that pass through q also go through u and v, and all contain copies of the edges (u, q) and (q, v). Let us fix a copy of (u, q) and call it e. In this case, the probability to move in $\mathcal{G}_P(C)$ from u to v through edge e in two steps is $1/(2\deg_{\mathcal{G}_P(C)}(u))$, since the probability to move along e is $1/\deg_{\mathcal{G}_P(C)}(u)$ and the probability to take a copy of (q, v) is $\frac{1}{2}$. If we combine these arguments with the fifth property specified in Lemma 18, we conclude that the probability to move from u to v along edge (u, v) in $\mathcal{G}_{P^*}(C^*)$ differs from the probability to do the corresponding move in $\mathcal{G}_P(C)$ by at most a factor of $O_\varepsilon(1)$. We further conclude inductively that if a random walk in $\mathcal{G}_{P^*}(C^*)$ moves in $t = O_\varepsilon(1)$ steps from vertex u to vertex v with probability p, then the same movement happens in $\mathcal{G}_P(C)$ with probability $\Omega_\varepsilon(p)$. Furthermore, by definition of well-contractible vertices, if the walk in $\mathcal{G}_{P^*}(C^*)$ contains a cycle of odd parity, then so does the corresponding walk in $\mathcal{G}_P(C)$.

It remains to address the probability of choosing u as the starting vertex. Since we may contract many vertices into u during our construction, the probability of choosing u as a starting vertex in $\mathcal{G}_{P^*}(C^*)$ can be significantly greater than the probability of choosing u in $\mathcal{G}_P(C)$. To this end, recall that the probability to choose u as a starting vertex is $p' = |P^{*-1}(u)|/|V|$. We note that with probability $\Omega_\varepsilon(p')$ we sample either u or a vertex $q \in Q$ that is contracted into u, i.e., a vertex q with $\gamma(q) = u$. By the definition of well-contractible vertices, the probability to move from q to u in the first step of the random walk is at least $\frac{1}{2}$. Summing up over all starting vertices (including u), we obtain that we end up at u (either in step 1 or 2 of the random walk) with probability at least $p'/2$. Thus, for every walk that happens in $\mathcal{G}_{P^*}(C^*)$ with probability $p \cdot p'$ there is a (set of) walks in $\mathcal{G}_P(C)$ such that one of them happens with probability $\Omega_\varepsilon(p \cdot p')$. Furthermore, if the walk in $\mathcal{G}_{P^*}(C^*)$ contains an odd-parity cycle then so does the walk in $\mathcal{G}_P(C)$. For different walks in $\mathcal{G}_{P^*}(C^*)$, the sets of corresponding walks in $\mathcal{G}_P(C)$ are disjoint. Thus, we do not double

count and our result follows by observing that the walk in $\mathcal{G}_P(C)$ has length at most $2t + 1 \leq 3t$. ■

6 Extending the analysis to minor-free graphs

While throughout the paper we focused on testing bipartiteness of planar graphs, our techniques can easily be extended to any class of minor-free graphs. Recall that a graph H is called a *minor* of a graph G if H can be obtained from G via a sequence of vertex and edge deletions and edge contractions. For any graph H, a graph G is called *H-minor-free* if H is not a minor of G. (For example, by Kuratowski's Theorem, a graph is planar if and only if it is $K_{3,3}$-minor-free and K_5-minor-free.)

Let us fix a graph H and consider the input graph G to be an H-minor-free graph. We now argue now that entire analysis presented in the previous sections easily extends to testing bipartiteness of G. The key observation is that our analysis in Sections 2–5 relies only on the following properties of planar graphs:

(i) the number of edges in a planar graph is $O(n)$, where n is the number of vertices (Fact 3),

(ii) every minor of a planar graph is planar (Fact 2),

(iii) a direct implication of the Klein-Plotkin-Rao theorem for planar graphs (Lemma 7).

The first two properties hold for any class of H-minor-free-graphs (that is, the second property would be that every minor of an H-minor-free-graph is H-minor-free). Since the Klein-Plotkin-Rao theorem holds for any minor-free graph as well (cf. [13]), so does a version of Lemma 7 with a slightly different constant hidden by the big O notation. Therefore, we can proceed with nearly identical analysis for H-minor-free graphs and arrive at the following version of Theorem 4.

Theorem 19 *Let H be a fixed graph. There are positive functions f and g such that for any H-minor-free-graph G:*

- *if G is bipartite, then* **Random-Bipartiteness-Exploration**(G, ε) *accepts G, and*
- *if G is ε-far from bipartite, then* **Random-Bipartiteness-Exploration**(G, ε) *rejects G with probability at least* 0.99.

7 Conclusions

In this paper we proved that bipartiteness is testable in constant time for arbitrary planar graphs. Our result was proven via a new type of analysis of random walks in planar graphs. Our analysis easily carries over to classes of graphs defined by arbitrary fixed forbidden minors.

This is merely the first step that poses the following main question:

What graph properties can be tested in constant time in minor-free graphs?

References

[1] N. Alon, E. Fischer, I. Newman, and A. Shapira. A combinatorial characterization of the testable graph properties: it's all about regularity. *SIAM Journal on Computing*, 39:143–167, 2009.

[2] N. Alon and M. Krivelevich. Testing k-colorability. *SIAM Journal on Discrete Mathematics*, 15(2):211-227, 2002.

[3] I. Benjamini, O. Schramm, and A. Shapira. Every minor-closed property of sparse graphs is testable. *Advances in Mathematics*, 223:2200–2218, 2010.

[4] A. Bogdanov, K. Obata, and L. Trevisan. A lower bound for testing 3-colorability in bounded-degree graphs. In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 93–102, 2002.

[5] A. Czumaj, A. Shapira, and C. Sohler. Testing hereditary properties of nonexpanding bounded-degree graphs. *SIAM Journal on Computing*, 38(6): 2499–2510, April 2009.

[6] O. Goldreich, editor. *Property Testing: Current Research and Surveys*. Lecture Notes in Computer Science 6390, Springer Verlag, Berlin, Heidelberg, December 2010.

[7] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4): 653–750, July 1998.

[8] O. Goldreich and D. Ron. Property testing in bounded degree graphs. *Algorithmica*, 32(2): 302–343, 2002.

[9] O. Goldreich and D. Ron. A sublinear bipartiteness tester for bounded degree graphs. *Combinatorica*, 19(3):335–373, 1999.

[10] O. Goldreich and D. Ron. On testing expansion in bounded-degree graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, Report No. 7, 2000.

[11] A. Hassidim, J. A. Kelner, H. N. Nguyen, and K. Onak. Local graph partitions for approximation and testing. In *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 22–31, 2009.

[12] T. Kaufman, M. Krivelevich, and D. Ron. Tight bounds for testing bipartiteness in general graphs. *SIAM Journal on Computing*, 33(6): 1441–1483, September 2004.

[13] P. Klein, S. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 682–690, 1993.

[14] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3): 261–277, 1988.

[15] S. Marko and D. Ron. Approximating the distance to properties in bounded-degree and general sparse graphs. *ACM Transactions on Algorithms*, 5(2), Article No. 22, March 2009.

[16] I. Newman and C. Sohler. Every property of hyperfinite graphs is testable. *SIAM Journal on Computing*, 42(3): 1095-1112, 2013.

[17] H. N. Nguyen and K. Onak. Constant-time approximation algorithms via local improvements. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 327–336, 2008.